

What is Claimed:

1. A method for providing a database view comprising transaction-consistent data reflecting the contents of a database at a specific point in time, said database comprising data elements and associated with a transaction log, said transaction log comprising active transactions and inactive transactions, said database view comprising difference storage for storing prior versions of at least one of said data elements from said database, said method comprising:
 - determining a split point on said transaction log corresponding to said point in time;
 - finding each transaction on said transaction log prior to said split point which performs modifications on said database;
 - storing each of said modifications in said difference storage;
 - finding each active transaction on said transaction log prior to said split point; and
 - undoing any corresponding modifications in said difference storage.
2. The method of claim 1, further comprising:
 - initializing said difference storage.
3. The method of claim 1, wherein said determining a split point on said transaction log corresponding to said point in time further comprises disabling log truncation.
4. The method of claim 1, wherein said database view further comprises a page table and wherein said method further comprises initializing said page table.
5. The method of claim 1, wherein said step of undoing any corresponding modifications in said difference storage comprises deleting said corresponding modifications.
6. The method of claim 1, wherein said step of undoing any corresponding modifications in said difference storage comprises:
 - reading corresponding unmodified data in said database; and
 - writing said corresponding unmodified data in said difference storage.
7. The method of claim 1, wherein each of said data elements comprises a page of data.

8. The method of claim 7, wherein said difference storage comprises at least one sparse file.
9. The method of claim 8, wherein said step of storing each of said modifications in said difference storage comprises allocating a region of memory in one of said sparse files.
10. The method of claim 9, where each of said data elements comprises a page of data, and where said database view further comprises a page table, and where said page table comprises, for each page:
 - first stored data indicating whether said page has been stored in said difference storage; and
 - second stored data indicating whether said region has been allocated in said difference storage.
11. The method of claim 10, where said method further comprises:
 - detecting that said page table is invalid;
 - for each region in said sparse files, determining whether said region has been allocated;
 - for each region in said sparse files, setting said second stored data based on whether said region has been allocated.
12. The method of claim 10, where a determination of whether data is stored in a specific page in said difference storage comprises:
 - checking said first stored data, and if said first stored data indicates that said specific page has been stored in said difference storage, determining that data is stored in said specific page in said difference storage;
 - checking said second stored data, and if said second stored data indicates that said region has not been allocated in said difference storage, determining that data is not stored in said specific page in said difference storage; and
 - if said first stored data does not indicate that said page has been stored in said difference storage and said second stored data does not indicate that said region has not been allocated in said difference storage, reading page data from a corresponding area of said difference storage for said specific page, and determining if said page data from said corresponding area is valid.

13. The method of claim 7, where said database view further comprises a page table, and where said page table comprises, for each page:

first stored data indicating whether said page has been stored in said difference storage.

14. The method of claim 13, where a determination of whether data is stored in a specific page in said difference storage comprises:

checking said first stored data, and if said first stored data indicates that said specific page has been stored in said difference storage, determining that data is stored in said specific page in said difference storage; and

if said first stored data does not indicate that said specific page has been stored in said difference storage, reading page data from a corresponding area of said difference storage for said specific page, and determining if said page data from said corresponding area is valid.

15. The method of claim 1, further comprising:

accepting a request for a specific data element in said database view;

determining if data is stored in a location corresponding to said specific data element in said difference storage;

responding to said request by reading said difference storage if data is stored in a location corresponding to said specific data element in said difference storage; and

responding to said request by reading said database if data is stored in a location corresponding to said specific data element in said difference storage.

16. The method of claim 15, where said step of determining if data is stored in a location corresponding to said specific data element in said difference storage comprises determining if said difference storage contains valid data in said location.

17. The method of claim 15, where said step of determining if data is stored in a location corresponding to said specific data element in said difference storage comprises consulting a page table.

18. The method of claim 1, where said method further comprises:

detecting a modification made to said database storing a first specific value to a location in said database in place of a second specific data element;

determining if a corresponding location in said database view contains valid data;
if said corresponding location in said database view does not contain valid data, writing said second specific data element in said corresponding location.

19. At least one of an operating system, a computer readable medium having stored thereon a plurality of computer-executable instructions, a co-processing device, a computing device, and a modulated data signal carrying computer executable instructions for performing the method of claim 1.

20. A system for providing a database view comprising transaction-consistent data reflecting the contents of a database at a specific point in time, said database comprising data elements and associated with a transaction log, said transaction log comprising active transactions and inactive transactions, said system comprising:

a split point determiner for determining a split point on said transaction log corresponding to said point in time;

a first transaction log analyzer for finding each transaction on said transaction log prior to said split point which performs modifications on said database;

difference storage for storing each of said modifications in said difference storage;

a second transaction log analyzer finding each active transaction on said transaction log prior to said split point; and

a difference storage modifier for undoing any corresponding modifications in said difference storage.

21. The system of claim 20, said split point determiner further comprising:
a log truncation disabler for disabling log truncation.

22. The system of claim 20, where said system further comprises:
a page table comprising data indicating whether a specific data element has been stored in said difference storage.

23. The system of claim 22, where said page table comprises:

first stored data indicating whether said specific data element has been stored in said difference storage.

24. The system of claim 23, where said difference storage comprises a sparse file and where said page table further comprises:

second stored data indicating whether a region corresponding to said specific data element has been allocated in said difference storage.

25. The system of claim 24, where said system further comprises:

a region allocation determiner for, for each region in said sparse files, determining whether said region has been allocated; and

a second stored data setter for, for each region in said sparse files, setting said second stored data based on whether said region has been allocated.

26. The system of claim 20, further comprising:

a request responder for accepting a request for a specific data element in said database view, determining if data is stored in a location corresponding to said specific data element in said difference storage, responding to said request by reading said difference storage if data is stored in a location corresponding to said specific data element in said difference storage; and responding to said request by reading said database if data is stored in a location corresponding to said specific data element in said difference storage.

27. A computer-readable medium for providing a database view comprising transaction-consistent data reflecting the contents of a database at a specific point in time, said database comprising data elements and associated with a transaction log, said transaction log comprising active transactions and inactive transactions, said database view comprising difference storage for storing prior versions of at least one of said data elements from said database, said computer-readable medium with instructions to perform acts comprising:

determining a split point on said transaction log corresponding to said point in time;

finding each transaction on said transaction log prior to said split point which performs modifications on said database;

storing each of said modifications in said difference storage;
finding each active transaction on said transaction log prior to said split point; and
undoing any corresponding modifications in said difference storage.

28. The computer-readable medium of claim 27, said acts further comprising:
initializing said difference storage.
29. The computer-readable medium of claim 27, said determining a split point on said transaction log corresponding to said point in time further comprising:
disabling log truncation.
30. The computer-readable medium of claim 27, where said database view further comprises a page table and where said acts further comprise:
initializing said page table.
31. The computer-readable medium of claim 27, where said step of undoing any corresponding modifications in said difference storage comprises:
deleting said corresponding modifications.
32. The computer-readable medium of claim 27, where said step of undoing any corresponding modifications in said difference storage comprises:
reading corresponding unmodified data in said database; and
writing said corresponding unmodified data in said difference storage.
33. The computer-readable medium of claim 27, where each of said data elements comprises a page of data.
34. The computer-readable medium of claim 33, where said difference storage comprises at least one sparse file.
35. The computer-readable medium of claim 34, where said act of storing each of said modifications in said difference storage comprises allocating a region of memory in one of said sparse files.

36. The computer-readable medium of claim 35, where each of said data elements comprises a page of data, and where said database view further comprises a page table, and where said page table comprises, for each page:

first stored data indicating whether said page has been stored in said difference storage; and
second stored data indicating whether said region has been allocated in said difference storage.

37. The computer-readable medium of claim 36, where said acts further comprise:

detecting that said page table is invalid;
for each region in said sparse files, determining whether said region has been allocated;
for each region in said sparse files, setting said second stored data based on whether said region has been allocated.

38. The computer-readable medium of claim 36, where a determination of whether data is stored in a specific page in said difference storage comprises:

checking said first stored data, and if said first stored data indicates that said specific page has been stored in said difference storage, determining that data is stored in said specific page in said difference storage;

checking said second stored data, and if said second stored data indicates that said region has not been allocated in said difference storage, determining that data is not stored in said specific page in said difference storage; and

if said first stored data does not indicate that said page has been stored in said difference storage and said second stored data does not indicate that said region has not been allocated in said difference storage, reading page data from a corresponding area of said difference storage for said specific page, and determining if said page data from said corresponding area is valid.

39. The computer-readable medium of claim 33, where said database view further comprises a page table, and where said page table comprises, for each page:

first stored data indicating whether said page has been stored in said difference storage.

40. The computer-readable medium of claim 39, where a determination of whether data is stored in a specific page in said difference storage comprises:

checking said first stored data, and if said first stored data indicates that said specific page has been stored in said difference storage, determining that data is stored in said specific page in said difference storage; and

if said first stored data does not indicate that said specific page has been stored in said difference storage, reading page data from a corresponding area of said difference storage for said specific page, and determining if said page data from said corresponding area is valid.

41. The computer-readable medium of claim 27, said acts further comprising:

accepting a request for a specific data element in said database view;

determining if data is stored in a location corresponding to said specific data element in said difference storage;

responding to said request by reading said difference storage if data is stored in a location corresponding to said specific data element in said difference storage; and

responding to said request by reading said database if data is stored in a location corresponding to said specific data element in said difference storage.

42. The computer-readable medium of claim 41, where said step of determining if data is stored in a location corresponding to said specific data element in said difference storage comprises determining if said difference storage contains valid data in said location.

43. The computer-readable medium of claim 41, where said act of determining if data is stored in a location corresponding to said specific data element in said difference storage comprises consulting a page table.

44. The computer-readable medium of claim 27, where said computer-readable medium further comprises:

detecting a modification made to said database storing a first specific value to a location in said database in place of a second specific data element;

determining if a corresponding location in said database view contains valid data;

MSFT-1746 / 301620.1

if said corresponding location in said database view does not contain valid data, writing said second specific data element in said corresponding location.